

Amendments to the Abstract:

Please amend the Abstract on page 65 as indicated:

ABSTRACT OF THE DISCLOSURE

METHOD AND APPARATUS FOR PERFORMING A STABLE HASH-BASED MAPPING COMPUTATION IN CONSTANT TIME OVER A DYNAMICALLY VARYING TARGET SET OF COMPUTATIONAL RESOURCES

A method is presented for mapping a source identifier in a source identifier space to a target identifier in a target identifier space using a hash-based computation that is stable over time with respect to a change in the number of target identifiers. A data item identifiable by a source identifier is to be associated with some type of computational resource that is represented by a target object identifiable by one or more target identifiers. The set of target objects is dynamically variable, yet the mapping is stable over time. After hashing the source identifier to produce an index position of an entry in a table, a target identifier is retrieved from the table entry, thereby mapping the source identifier to the target identifier in a mapping operation whose speed is independent of the number of target identifiers. Each entry in the table is related to a single target identifier, yet each target identifier may be related to more than one table entry. The target becomes related to a table entry based on a ["]nearness["] computation that depends upon the table index position of the table entry and a target identifier for the related target. The nearness computation is performed between each table entry and each target identifier to obtain a fair distribution of relationships between table entries and targets. Targets can be added or removed with minimal impact on the table. The mapping operation may also incorporate target weighting that is proportional to a target's computational capacity.